

AUTHOR:

Dr Michael Pearson
Principal Research Fellow
Employment Research Institute
Edinburgh Napier University
Colinton Road
Edinburgh EH14 1DJ
Tel. +44 – 131 4554361
E-mail:
drmpearson@hotmail.co.uk
m.pearson@napier.ac.uk

The Shape of Uncertainty in Revenue Management

The Shape of Uncertainty in Revenue Management

Abstract

The revenue management problem we address is traditionally posed by the budget airlines. These airlines do not tend to have protected classes but instead use dynamic pricing to determine allocation policy. This usually means that passengers who book early get discounted rates, while those booking closer to the date of the flight pay higher rates. We expand on earlier work with the newsvendor problem to employ a network equilibrium solution which takes account of both customer and airline objectives to formulate a dynamic solution to the pricing and allocation of airline tickets in an uncertain demand and supply market. We make use of phase diagram analysis when identifying performance during two stages of the pricing strategy consistent with the airline demand and pricing information. One phase diagram (endogenous) identifies efficient frontiers jointly associated with customer satisfaction and airline performance. The other (exogenous) identifies optimal performance with regard to the global market place. The optimal solution strategy is achieved by a stochastic differential equation set with price and booking closure as control variables. The solution is illustrated with data from the airline industry. The interaction between goodwill costs and the price elasticity of demand for varying levels of forecasting accuracy shows that contribution to profit benefits from i) the incorporation of goodwill costs ii) low levels of elasticity for good forecasting iii) high levels of elasticity for poor forecasting. We identify a breakeven level of forecasting stability which establishes an equilibrium between lean and agile practice.

Key Words: Revenue Management, Dynamic Pricing, Airline Booking, Stochastic Scheduling, Customer Goodwill

1. Introduction

Yield and revenue management are becoming increasingly important following on from success in the airline and hotel business to the widespread use of online sales and booking systems. Car rental, restaurants, cargo shipping, entertainment, and the manufacturing and fashion industries are all new areas of implementation. Most methods incorporate the finite supply of an item of inventory during a limited time horizon for sales under conditions of varying and often uncertain demand frequently defined in product-oriented terms which fail to take into account customer behaviour. Assumptions are often made about the demand and supply distributions, for example that they follow a Poisson, normal or other distribution. This can lead to a limitation in the range of application of the methodology to areas where such distributions occur. Another limitation is caused by ignoring the ability of an operator to overcome fluctuations in demand and supply through good forecasting and

demand modelling methods. There is a need for a broader analysis which incorporates the effect of uncertainty, measured using prediction or forecasting errors, on airline profit.

Unpredictable variability results in uncertainty, while uncertainty becomes risk when it is manageably quantifiable. In order to map prediction variability, risk and uncertainty we need an applicable mapping structure or topology which is able to illustrate different dimensions of the business model. Changes in forecast accuracy might be attributable to features of the market such as the degree of competitiveness and uncertainty experienced when entering new trading markets. Some businesses might be content to explore smaller scale niche markets for their products where the uncertainty is lower and perhaps more quantifiable. Others may seek to expand their business by entering more profitable high volume global trading markets with greater attached risk to the venture (Figure 1a). What both types of business would want to do is optimise their expected profit levels within the trading environment of uncertainty.

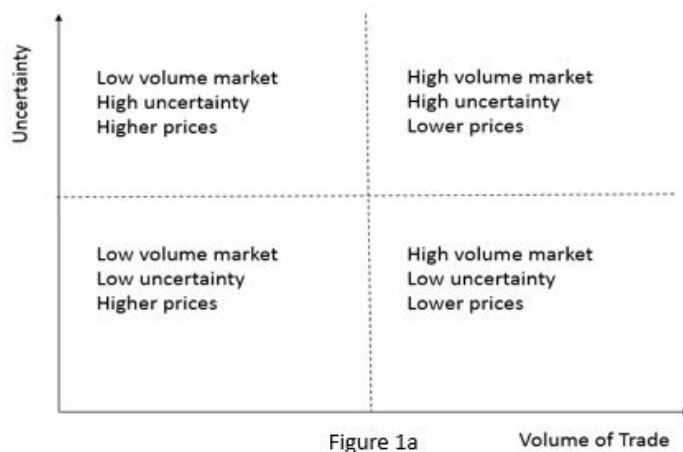


Figure 1a

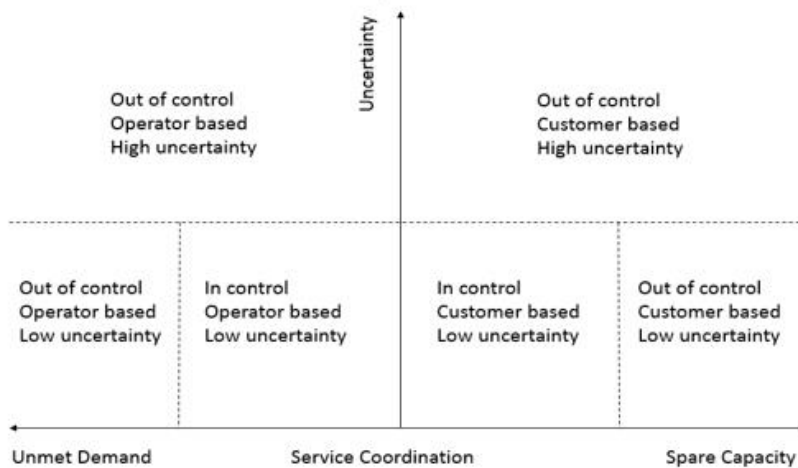


Figure 1b

Figures 1a and 1b. Global and Local Topologies

Another dimension of the graphical layout is the level of uncertainty arising between the customer and the business operator. An operator might be prepared to sacrifice customer goodwill for

the ability to cancel services at short notice in order to run a more cost efficient business.

Disappointed customers might decide to change future plans and even turn to another operator which creates uncertainty for both the business operator and the customer at the local level, identified as the interface between customer and business operator (Figure 1b). We will argue that the mapping structures of these two levels, global and local, are distinct, with the global topology being an open one with a potentially unbounded efficient frontier, while the local topology is closed, with bounded optimality being achieved at the intersection of three efficient frontiers.

Figures 1a and 1b illustrate a way of mapping uncertainty in a business framework where transactions occur in a global (operator trading) and local (customer/operator) context. The objective in both cases is to employ an optimal business strategy depending on the type of market entered as well as the customer base being sought. Figure 1a generates four segments depending on the uncertainty level and the size or volume of the market trade. Operators entering the high volume market frequently meet with stiff competition and volume discounts as well as greater market uncertainty. On the other hand finding a niche market may result in low uncertainty, low volume output with higher prices. Figure 1b segments the uncertainty according to a customer/operator based service coordination policy. The customer tends to benefit from the availability of spare capacity experienced as maintaining goodwill although this risks the operator's profit margins. On the other hand the operator tends to benefit from increased demand although at the risk of losing future custom. Neither the operator nor the customer benefit from high uncertainty levels and so the preferred segments tend to be those with low uncertainty and without extremes of spare capacity or unmet demand, deemed to be 'in control'. The remaining challenge is mapping the optimal business strategy, such as maximising contribution to profit, to the identified topology of uncertainty.

In businesses employing revenue management the effects of forecasting errors can be incorporated into customer service and goodwill as key aspects of the way customers make choices and operators run their businesses. In the airline business, not being able to find a suitably priced seat can affect a customer's propensity to seek future bookings from that airline and, although this goodwill cost is difficult to assess, it is considered particularly acute in competitive markets and with customers who are frequent air travellers. Goodwill costs may also be associated with the *customer's* personal inconvenience at having to find a booking elsewhere in addition to the *airline's* loss. Such costs therefore refer to a longer term feature such as a sense of loyalty and the desire for customers to return to the airline in the future having experienced a disappointment while making a change to the itinerary such as a name change or flight adjustment which may result in a cancelled booking and also deter the customer from making future bookings.

The incorporation of the price elasticity of different demand segments into a stochastic revenue management model reflects customer sensitivity so that airlines can set prices accordingly for each segment to maximise total revenue as well as maintain goodwill. The segments themselves should have different quality of service characteristics leading to product discrimination while within each

segment price discrimination might also be used depending on booking progress. Models have been developed which have sophisticated price response functions but rarely in combination with other factors such as uncertainty (demand risk), goodwill and the distinction between global (exogenous) and local (endogenous) influences. There is a need for an effective tool to monitor and identify patterns of uncertainty in revenue management in order to understand and reduce the effects of variability in the global market place as well as due to local transactions between customer and business.

The quality of forecasting is generally measured using forecasting errors. A good forecast is made when the errors are small so that the forecast mirrors the actual demand. The poorest forecasting occurs when nothing is inferred about future demand, which is then modelled using naïve methods. That is the forecast of future demand is modelled as the demand for the previous time point. In order to identify patterns of uncertainty we simulate forecasting errors using a structured experimental design. At one extreme we use good forecasting, which is the demand with a small random change added. At the other end we have poor, that is naïve, forecasting. Between these two extremes we simulate a weighted average of the two methods reflecting the quality of forecasting. In addition to the quality of forecasting, the experimental design includes two other factors: the level of goodwill and the price elasticity of demand. The experiment is then repeated at several factor levels and the shape of uncertainty reflected by the phase diagrams is analysed.

The remainder of the paper is organised as follows. The literature on revenue and yield management and dynamic pricing is reviewed in Section 2. The methods, including problem outline, terminology, formulation and solution are outlined in Section 3. In Section 4 we introduce a case study using airline data and use it to illustrate the methodology as well as carry out sensitivity analysis. The paper ends with a conclusion in Section 5.

2. Literature Review

There is extensive literature on revenue and yield management including surveys published by Weatherford and Bodily (1992), McGill and van Ryzin (1999), Elmaghraby and Keskinocak (2003), as well as books by Talluri and van Ryzin (2004), Phillips (2005) and Belobaba et al. (2009). Levin et al. (2009) introduce a special issue on revenue management and dynamic pricing with contributions spanning a wide range of applications. The earlier yield management solutions proposed by Littlewood (1972) and Belobaba (1989) require some knowledge of the demand distribution, such as the demand densities of future requests, in order to calculate appropriate seat allocation values. Levin et al. (2008) describe a risk-adjusted method which incorporates a simple risk measure permitting control of the probability that total revenues fall below a minimum acceptable level under the assumption that demand follows a price-dependent, non-homogeneous Poisson process. Some methods maximise expected revenue either by offering multiple product classes at different prices or by the use of a single product class with dynamically varying price over time (Levin et al., 2008).

Related newsvendor solutions have been proposed which are distribution-free. One of the earliest was Scarf (1958) based on setting upper and lower bounds on the optimal solution across all distributions. This approach extends the generality of the methodology at the expense of failing to identify an exact solution often required by management allocation systems. Later work has extended this methodology while also repeating the lower and upper bound approach (Alfares et al., 2005; Raza, 2014). In this article we allow for a distribution free approach while restricting the forecasting errors to a normal distribution in order to identify an exact solution.

Some modelling takes into account customer behaviour factors (Ming et al., 2009; Dana, 2008) and demand-based pricing (Belobaba et al., 2009) as well as the external environment (Ozkan et al., 2014). A key concept in measuring the effect of pricing policy on consumer welfare is consumer surplus (Phillips, 2005), which is the difference between a customer's willingness to pay and the price at which he purchases (assumed zero if he does not purchase). The price elasticity of demand is incorporated into demand-based pricing (Chap. 4, Belobaba et al., 2009) on the assumption that some customers are willing to pay a high price for the convenience of air travel, while others will only fly at lower prices. A modelling variation assumes that customers act strategically and bid for units of a fixed capacity over time (Vulcano et al., 2002). We take into account both behavioural factors, such as the effect of consumer surplus, as well as the external environment associated with market and pricing trends.

Goodwill costs can also be taken into account when identifying an optimal allocation policy. An approximate assessment of the airline's cost of a permanently lost customer is the expected net present value of all future bookings from the customer minus the opportunity costs of those bookings (McGill and Van Ryzin, 1999). These costs are frequently ignored in the allocation of airline passenger seats, although Alfares et al. (2005) did include the incorporation of goodwill or shortage costs to update Scarf's distribution-free algorithm (Scarf, 1958). Goodwill costs might also reflect the partial refund of fares, which, it is argued, is optimal in both revenue and social terms (Gallego et al., 2010). We make use of a 'mix' phase diagram specifically to monitor the effect on passenger goodwill.

The literature on forecast accuracy in revenue management is limited by context and generality. Lee (1990) concluded that: (a) the mean of the forecast demand was more important than the standard deviation; (b) under-forecasting has less impact on revenues than over-forecasting; (c) expected revenues rise by 0.5-3% for each 10% rise in the accuracy of forecasting. Poelt (1998) also translated a reduction in forecast errors to an increase in revenue, while Weatherford et al. (2002) found that the accuracy of demand forecasts was of the greatest importance. Birbil et al. (2009) introduce robust versions of the single leg problem to take into account inaccurate estimates of the underlying demand distributions. The rationale behind their method is for management to immunize the revenues against demand uncertainty. The method uses an adapted demand probability distribution

to generate a solution which improves not only the average but also the standard deviation of the solution revenues.

In another approach to reducing the effect of uncertainty, Marcus et al. (2008) model stochastic dynamic pricing employed by a budget operator in competition with a full-service carrier and the effect that randomness has on optimal under-pricing behaviour by introducing oscillating behaviour patterns. Sahin et al. (2009) quantify the economic impact of having uncertainty on the inventory level in a newsvendor context. Modarres et al. (2009) analyse the effect of uncertainty in capacity allocation. However, the variability of the forecasting errors is not analysed and mapped directly by any of these approaches, though issues of bias are addressed as well as the effect of the variability of actual fare values and the effects on revenues. The method we adopt takes into account the effect of the variability of the forecasting errors on a range of factors, including customer goodwill, price elasticity as well as the airline contribution to profit.

3. Methods

3.1. Problem Outline

Earlier research (Kocabiyıkoğlu et. al., 2014) has shown that, when coordinating decisions on prices and capacity allocation using nesting, near-optimal performance is usually achieved by a simple hierarchical policy that sets prices first, based on a non-nested stochastic model, and then uses these prices to optimize nested capacity allocation. Furthermore the largest revenue benefits stem from adjusting prices to account for demand risk. One response to this might be to abandon the management of capacity allocation altogether since a fare class can be closed by setting sufficiently high prices. Such a full implementation of dynamic pricing rarely happens in practice (Kocabiyıkoğlu et. al., 2014), however, and we will use a method whereby decisions on prices and capacity allocation are coordinated taking into account the risk factors.

The problem we study, then, simplifies the stochastic dynamic programming approach in order to illustrate the structure of demand uncertainty arising from different forecasting methods. We assume that there are two classes of passengers, low fare passengers who generally arrive before high-fare passengers and that there is a two tier pricing strategy as a starting point based on management knowledge as well as historical data. We use a demand profile with 23 time intervals from booking information extracted from airline data. The first 12 time intervals define the first stage with associated low price demand profile. The remaining 11 time intervals define the second stage associated with high ticket price. We adjust the starting price at each time interval based on the remaining available capacity. If there is a shortage of such capacity then the price is increased and decreased otherwise. A transformation is used to distinguish between a local “mix” variable which highlights the coordination between the customer and the airline operator and the maintenance of goodwill, and a global variable which reflects competition in the market place modelled using the price elasticity of demand.

We assume that that the forecasting errors are normally distributed to allow for the development of path trajectories in significantly uncorrelated phase diagrams, which plot forecast error variability against the local and global variables, and monitor the transition from controlled risk to high uncertainty. The methodology makes use of four efficient frontiers designed to map and monitor local and global uncertainty in the context of optimal contribution to profit. Beginning with stage one, a capacitated stochastic dynamic pricing algorithm sequentially adjusts the available seats and price at each time interval in order to maximise, subject to capacity limitations, the profit function, which incorporates both airline and passenger revenue and costs, since previous research has shown that there is only a channel equilibrium solution for the stochastic problem (Pasternack, 1985; Pearson, 2003).

We address the single-leg yield management problem through the analysis of the forecasting errors of the demand rather than the variability of the demand. The reason is that a good forecaster (demand manager) will produce better forecasts (predictions) with variable demand data than a poor forecaster. It is not, therefore, fluctuations in demand but the forecaster's ability to predict demand that determines good performance. We also assume that the forecaster is doing a good job, resulting in an unbiased forecast before testing our model beyond this assumption. Finally, although we perform simulations using forecasting quality, goodwill costs and elasticity of demand as control variables in the objective function, we measure outcomes from the point of view of the airline using the airline ticket revenue as dependent variable when plotting these outcomes. In this way we hope to avoid the criticism that goodwill costs are difficult to quantify but instead demonstrate that incorporating such costs in the objective function benefits airline profit. It is then for the airline operator to discern, through customer feedback and research, which efforts at improving goodwill are effectively modelled by control levels in the objective function.

3.2 Terminology

Table 1 introduces the symbols used during the analysis. Some of these, such as the terms 'overage' and 'underage' may need illustrating. Traditionally these terms are associated with the newsboy problem, whereby a newsboy takes copies of a newspaper to sell on the street. If he takes too many then he has some surplus copies which are his overage. If he takes too few he has a shortage of copies known as underage. So, if supply is greater than demand the newsboy has overage while if supply is less than demand he has underage. We apply these terms to the airline operator selling passenger seats. If there are spare seats these are the airline operator's overage as the supply of seats exceeds the demand for them. Similarly, if there are no seats left while there remains some demand then the difference between the demand and the supply is the airline operator's underage.

We extend this usage by referring to the customer's 'overage', which is when he books a seat but does not use it. Similarly the customer's underage occurs when he is unable to book a seat in his

price bracket because there is not one available (usually due to the airline closing the booking in this price range).

Symbol	Interpretation
c_{u_1}	<i>Goodwill Cost of customer disappointment at unavailable seat or (this cost – Compensation) or weighted average</i>
c_{o_2}	<i>Goodwill Cost to airline of concern at loss of customer (Overbooking costs can be incorporated if necessary)</i>
c_{p_0}	<i>Baseline contribution to profit (Baseline ticket cost: determined by the historical profile and the market). This is fixed at each stage</i>
c_{o_1}	<i>Cost to customer of booking and not using seat (such as Ticket Cost – Refund) averaged over all bookings</i>
c_{u_2}	<i>Cost to airline of unsold seats (such as Ticket Cost or Refund or weighted average)</i>
c_p	<i>Variable contribution to profit (Variable ticket cost : used in optimisation methods)</i>
k	<i>Safety factor (determines amount of additional seats, $k\sigma_e$, to hedge against forecasting uncertainty)</i>
$\Phi(k)$	<i>Cumulative Normal Distribution Function (Steady state newsvendor ratio)</i>
$\phi(k)$	<i>Normal Density Function</i>
μ_D	<i>Expected demand</i>
$\widehat{\mu}_D$	<i>Forecast demand</i>
μ_Q	<i>Expected supply of seats</i>
σ_e	<i>Standard deviation of forecasting errors</i>
C	<i>Seat capacity (updated during allocation process)</i>
<i>%Underage</i>	<i>Underage as a % of mean demand</i>
<i>%Overage</i>	<i>Overage as a % of mean demand</i>
<i>Error CV</i>	<i>Error standard deviation divided by mean demand</i>

Table 1. Definition of terms

The marginal contribution to profit c_p is also the cost of the airline ticket since overheads are assumed constant. The overage and underage costs of the downstream operator (this is usually the customer buying the seat, but it could also be an intermediary buying seats for other users) are c_{o_1} and c_{u_1} respectively, while for the upstream operator (usually the airline operating the flights) they are c_{o_2} and

c_{u_2} . An interesting feature of the problem and the way it is formulated is that the downstream operator's overage is the same as the upstream operator's underage, though they may have different attitudes to these phenomena resulting in unequal costs. For instance a customer, whose decision variable is the number of tickets, Q , to order, may buy two tickets one for himself and one for a business colleague and, in the end, only need one. This is overage from the point of view of the customer since $Q > D$. The customer may set the price of overage c_{o_1} at the cost of the ticket if no refund is possible under the airline terms and conditions. If a refund is possible then the cost of overage is set at the difference between the full price paid and the refund price. From the alternative point of view of the airline operator, whose decision variable is the demand, D , generated through pricing and other restrictions, there is now an empty seat which is less than a full airline resulting in underage, costing c_{u_2} , since $D < Q$. The cost to the airline of this underage may be set at the refund price, assuming that the seat is not sold in that time period. If the seat fare is neither refunded nor levied because of low demand then the cost of underage is simply equal to the full fare. There are other aspects of the problem such as overbooking, which occurs close to scheduled flight take-off and requires that the airline persuade some passengers to take another flight.

3.3 Problem Formulation

The revenue management problem we address is posed by both budget and traditional airlines which employ demand-based pricing methods based on dynamic pricing to determine allocation policy (Phillips, 2005). This usually means that passengers who book early get discounted rates, while those booking closer to the date of the flight pay higher rates, though this can change depending on the state of supply and demand.




Figure 2. Remaining capacity (C) seat demand (μ_D) and supply allocation (μ_Q)

Figure 2 illustrates the way in which we model the airline yield management problem. The allocation decision is made by setting the value of the safety factor, k , at time t . This determines the amount of seats, $k\sigma_e$ in addition to expected demand to hedge against forecasting uncertainty so that the objective function is optimised. The decision is made repeatedly during the period from the opening of the booking schedule to flight take-off. During this period the values of the expected demand, μ_D expected supply, μ_Q the safety factor, k , the standard deviation of the forecasting errors, σ_e , and the remaining capacity, C , are changing as the bookings progress. So also is the value of the seat fare which is

adjusted to ensure that sufficient capacity remains after each time period. Lowering the price results in increased demand in the next time period which restores excessive remaining capacity while increasing the price has the reverse effect. We note that the cabin capacity is not necessarily equal to the supply allocation. If the capacity is less than the allocation then overbooking occurs. If it is greater than the allocation then seats may be protected for allocation to high fare customers at the time of take-off to ensure customer satisfaction by avoiding disappointment. In our formulation we avoid overbooking by limiting the number of seat bookings to the capacity of the airplane.

3.4 Transformation of Variables

We transform the variables (Pearson, 2003; 2010) for the expected demand (μ_D) and supply (μ_Q) by use of a sum difference transformation so that $\mu = \mu_Q - \mu_D$ and $\eta = \mu_Q + \mu_D$. The transformed variable, μ , is the local variable reflecting the way in which the airline relates to its customers (or other partners in the transaction; see Figure 1b), while the transformed variable, η , is the global variable reflecting pressure from outside competition and market forces (see Figure 1a). In our analysis the value of the expected demand is determined from airline data. The value of the expected seat allocation (supply) is found by solving a stochastic differential equation, set up using an objective function, together with dynamic pricing methods based on elasticity of demand differentiated between different types of passengers, described later in this article.

3.5 Assumptions

Our model has six assumptions:

- 1.) That the overhead costs are fixed so that a change in the ticket price is accurately reflected by a change in the contribution to profit.
- 2.) Unbiased demand and supply fitting or forecast techniques are applied.
- 3.) The relationships between the standard deviation of the errors and the demand and supply means are well behaved.
- 4.) That the global (η), but not the mix (μ), variable affects price changes so that c_p does not depend on μ since $\eta = \mu_D + \mu_Q = 2\mu_D + \mu$. This is an approximation based on the fact that changes in the mix variable, μ , occur in order to improve smooth coordination between the customer and the airline operator maintaining goodwill, while changes in the global variable reflect shifts to high volume trading and competition in the market place modelled using the price elasticity of demand.
- 5.) The prediction errors are normally distributed
- 6.) That price elasticity of demand for airline seat tickets applies, so that increasing the price of a ticket decreases demand and vice versa.

3.6 Objective Function

The objective is to maximise expected profit:

$$E\{\text{Profit}\} = E\{\text{Contribution from captured demand} - \text{Costs of overage} - \text{Costs of underage}\}$$

$$\text{Now } E\{\text{Contribution from captured demand}\} = \mu_D c_p - \{(\phi(k) - k(1 - \Phi(k)))c_p\} \sigma_e$$

$$\text{And } E\{\text{Cost of overage}\} = \{(k\Phi(k) + \phi(k))(c_{o_1} + c_{u_2})\} \sigma_e$$

$$\text{While } E\{\text{Cost of underage}\} = \{(\phi(k) - k(1 - \Phi(k)))(c_{u_1} + c_{o_2})\} \sigma_e$$

So that

$$E\{\text{Profit}\} = \mu_D c_p - \{(k\Phi(k) + \phi(k))(c_{o_1} + c_{u_2}) + (\phi(k) - k(1 - \Phi(k)))(c_{u_1} + c_{o_2} + c_p)\} \sigma_e \quad (1)$$

$$\text{subject to: } \mu - k\sigma_e = 0, \text{ (Newsvendor Constraint)} \quad (2)$$

where $\phi(k)$ and $\Phi(k)$ are the normal distribution density and cumulative distribution functions, respectively, for safety factor, k and the forecasting errors are assumed to have a normal distribution.

We are then able to find the maximum value of the objective function (1) subject to the constraint (2) by applying a Lagrange multiplier with decision variables μ , η and k . We then incorporate price elasticity assuming that the ticket price depends on η but not μ .

3.7 The equilibrium solution under changing variability

The equilibrium solution under changing variability is derived in Pearson (2003) and requires that both of the following equations hold (a summary of the proof is in Appendix A):

$$\Phi(k) + \frac{\partial \sigma_e}{\partial \mu} \phi(k) = \frac{c_{u_1} + c_{o_2} + c_p / 2}{c_{o_1} + c_{u_1} + c_{o_2} + c_{u_2} + c_p} \quad (3)$$

$$\phi(k) \frac{\partial \sigma_e}{\partial \eta} = \frac{c_p}{2(c_{o_1} + c_{u_1} + c_{o_2} + c_{u_2} + c_p)} \quad (4)$$

Eq. (3) is the ‘mix’ (overage/underage) solution, which tracks the way partners across decision frontiers synchronize their efforts to reach optimality (Figure 1b), and Eq. (4) is the ‘global’ (volume) solution, which tracks the relationship between uncertainty and the volume of trade experienced in the market sector (Figure 1a). Together they form a dynamic system of stochastic differential equations, which trace the optimal solution in circumstances where uncertainty increases or decreases over time and also with relation to differing contractual and marketing strategies.

3.8 The steady State Solution

A feature of the problem is that when $C = \mu_Q$ the optimal solution is

$$\Phi(k) = \frac{c_{u_1} + c_{o_2} + c_p}{c_{o_1} + c_{u_1} + c_{o_2} + c_{u_2} + c_p} \quad (5)$$

We will call this the *steady state solution* since it corresponds to the classical newsvendor problem with constant rate of variability. It is found by substituting $\mu_Q = C$ so that $\mu = C - \mu_D$ and

$$\frac{\partial \sigma_e}{\partial \mu} = \frac{\partial \sigma_e}{\partial (C - \mu_D)} = -\frac{\partial \sigma_e}{\partial \mu_D}$$

while $\eta = C + \mu_D$ and $\frac{\partial \sigma_e}{\partial \eta} = \frac{\partial \sigma_e}{\partial (C + \mu_D)} = \frac{\partial \sigma_e}{\partial \mu_D}$. If we substitute these results into Eqs. (3) and (4)

we get Eq. (5). We interpret this result by noting that the ability to choose a value $k = (\mu_Q - \mu_D) / \sigma_e$ such that $\mu_Q = C$ requires that the error variability (σ_e) does not depend on the expected supply (μ_Q), which is now a constant. The solution therefore reverts to the steady state solution, which serves as a valuable reference point. It also means that the rate of change of the error variability with respect to the expected demand is found from Eq. (4) to be

$$\frac{\partial \sigma_e}{\partial \mu_D} = \frac{c_p}{2\phi(k)(c_{o_1} + c_{u_1} + c_{o_2} + c_{u_2} + c_p)} \quad (6)$$

3.9 Equilibrium solution incorporating price elasticity

We incorporate price elasticity within our equilibrium solution by making use of the assumptions outlined in Section 3.5. The fifth assumption, that the forecasting errors are normally distributed, is the result of skilled forecasting, though there may be instances where the assumption may break down. One such instance is when the demand is very low which may result in a Poisson distribution. The equilibrium solution under conditions of changing variability can be shown, by using a Lagrange multiplier (see Appendix A), to be

$$\Phi(k) + \frac{\partial \sigma_e}{\partial \mu} \phi(k) = \frac{c_{u_1} + c_{o_2} + c_p / 2}{c_{o_1} + c_{u_1} + c_{o_2} + c_{u_2} + c_p} \quad (8)$$

$$\phi(k) \frac{\partial \sigma_e}{\partial \eta} = \frac{c_p / 2 + [\mu_D - (\phi(k) - k(1 - \Phi(k)))\sigma_e] \partial c_p / \partial \eta}{c_{o_1} + c_{u_1} + c_{o_2} + c_{u_2} + c_p} \quad (9)$$

As in Eq. (5) a feature of the problem is that when $C = \mu_Q$ we again make use of the fact that

$$\frac{\partial \mu_D}{\partial \eta} = -\frac{\partial \mu_D}{\partial \mu} \text{ when we add Eqs. (8) and (9) to get the optimal solution. Now we have to take into}$$

account the dependence of the forecast variability, σ_e on the contribution to profit, c_p as well as the dependence of the contribution to profit on the global variable, η . We find that

$$\Phi(k) + \frac{\partial \sigma_e}{\partial c_p} \frac{\partial c_p}{\partial \eta} = \frac{c_{u_1} + c_{o_2} + c_p + [\mu_D - (\phi(k) - k(1 - \Phi(k)))\sigma_e] \partial c_p / \partial \eta}{c_{o_1} + c_{u_1} + c_{o_2} + c_{u_2} + c_p} \quad (10)$$

since $\frac{\partial \sigma_e}{\partial \eta} = \frac{\partial \sigma_e}{\partial c_p} \frac{\partial c_p}{\partial \eta} + \frac{\partial \sigma_e}{\partial \mu_D} \frac{\partial \mu_D}{\partial \eta}$ and $\frac{\partial \sigma_e}{\partial \mu} = \frac{\partial \sigma_e}{\partial c_p} \frac{\partial c_p}{\partial \mu} + \frac{\partial \sigma_e}{\partial \mu_D} \frac{\partial \mu_D}{\partial \mu}$ and it is assumed that c_p does not depend on μ so that $\frac{\partial c_p}{\partial \mu} = 0$. Eq. (10) now forms the basis of our subsequent analysis. Two features are worthy of comment. Firstly the rate of change of the contribution to profit (c_p , reflected by the ticket price) with respect to volume (η , reflected by the ticket sales) is typically negative (assumptions 4 and 6). Secondly the rate of change of the variability of the forecasting errors with respect to the marginal contribution to profit can be modelled by considering current practice pricing policy.

3.10 Assumed Link between Error Variability and Current Practice Pricing

We model forecasting error variability as being at its lowest in the area of current practice pricing policy. In this way we reduce the risk of making bland assumptions and predictions about demand for seats in an area of pricing remote from current practice. Figure 3 illustrates this where c_{p_0} is the initial current practice price and the equation linking error variability to a revised price, c_p is given by:

$$\sigma_e = \sigma_{e_0} \left\{ r_\sigma \left(\frac{c_p}{c_{p_0}} - 1 \right)^2 + 1 \right\} \quad (11)$$

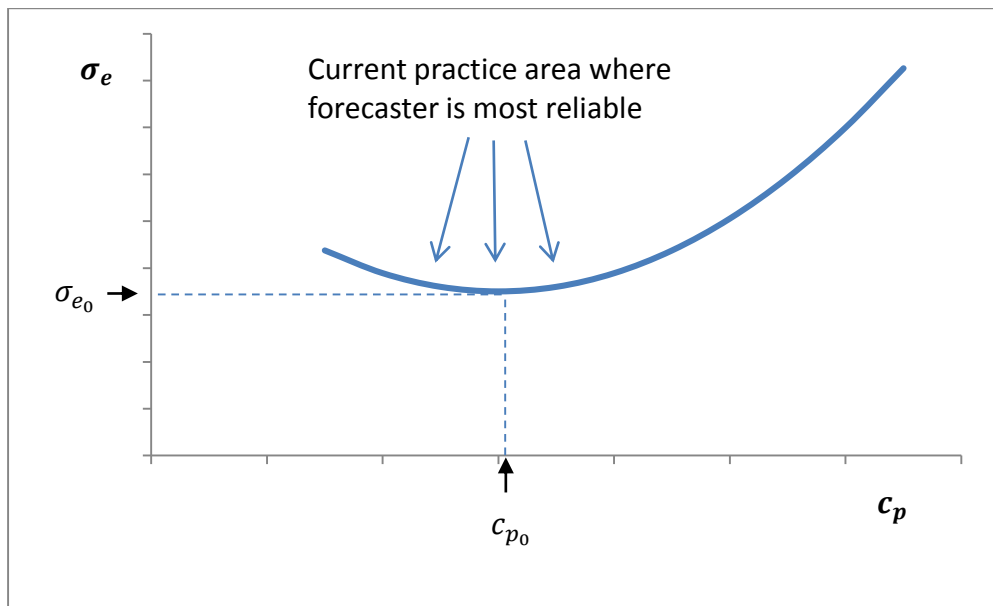


Figure 3. Error Variability versus Contribution to Profit (Price)

We note that r_σ controls the rate at which the quadratic expression for variability changes with respect to price. Therefore changing the price from the current practice level, c_{p_0} typically has the effect of increasing the forecast variability since the forecasting practitioner is operating outside of his familiar territory. If the forecaster has knowledge of the broader market with respect to pricing then a lower (or even zero) value of r_σ can be used. In our modelling we set $r_\sigma = 0.1$ so that doubling the current

practice price is associated with a 10% increase in the error standard deviation while trebling the price is associated with a 40% increase. We note that the rate of change of error variability with respect to the contribution to profit is:

$$\frac{\partial \sigma_e}{\partial c_p} = 2\sigma_{e_0} r_\sigma \left(\frac{c_p}{c_{p_0}} - 1 \right) / c_{p_0} \quad (12)$$

We use this in Eq. (10), together with the approximation $\frac{\partial \mu_D}{\partial c_p} \left(= \frac{1}{2} \frac{\partial \eta}{\partial c_p} \right)$ based on assumption 4, to calculate the optimum value of k based on the price elasticity.

3.11 Capacitating the Solution by Incremental Price Adjustments using Price Elasticity

The equilibrium solution derived in Equation 10 requires that $C = \mu_D$. We therefore employ a mechanism to control for this outcome by adjusting the ticket price appropriately. If there is a surplus of remaining capacity then the ticket price falls to stimulate demand to occupy the surplus. If there is a capacity shortfall then the ticket price is increased. The amount by which the ticket price is adjusted is determined using the price elasticity of demand. Price elasticity is the percentage increase in demand for a one percent increase in price, so that:

$$\text{Elasticity} = \frac{\frac{d\mu_D}{\mu_D}}{\frac{dc_p}{c_p}} = \frac{c_p}{\mu_D} \times \frac{d\mu_D}{dc_p} = \frac{c_p}{\mu_D} \times \frac{1}{2} \times \frac{d\eta}{dc_p} \quad (13)$$

We use this to calculate the forecast demand at a certain ticket price based on the forecast demand at the base price, c_{p_0} as well as the amount to increase the ticket price in the event of a shortage in remaining capacity. The formulae we use are:

$$\begin{aligned} \text{Fcst Demand at Price, } c_p &= \text{Fcst Demand at Price, } c_{p_0} + \frac{c_p - c_{p_0}}{\frac{dc_p}{d\mu_D}} \\ &= \mu_D + (c_p - c_{p_0}) \times \frac{\mu_D}{c_p} \times \text{Elasticity} \end{aligned} \quad (14a)$$

$$\text{Increase in ticket price} = \frac{\text{capacity shortage} - c_{p_0}}{\text{Elasticity} \times \mu_D - \text{capacity shortage}} \quad (14b)$$

where the shortage in remaining capacity is the difference between the actual remaining capacity and the remaining capacity forecast at the end of the previous time period. If there is a surplus of remaining capacity then the ticket price will fall. Equation 14b requires that the denominator is non zero.

3.12 Description of the Algorithm

Adjustments to the ticket price are therefore used to maintain the remaining capacity requirements, while the optimal allocation (value of k calculated from Equation 10) determines when to close the booking for that time period. We describe the algorithm below:

- i) Input demand profile for periods 1-23
- ii) Input ticket, underage and price elasticity values and targets for each stage
- iii) Check that demand profile meets capacity requirements and adjust accordingly

- iv) Set $j=0$
- v) Set $i = 0$
- vi) $j = j + 1$
- vii) Calculate steady state safety factor for stage j
- viii) $i = i + 1$
- ix) Generate forecast demand for time period i using naïve or random methods
- x) Calculate the standard deviation of the forecasting errors up to the current time period
- xi) Calculate remaining capacity shortage/surplus for previous time period
- xii) Calculate the ticket price adjusted for remaining capacity from Equation 14b
- xiii) Calculate the forecast demand using updated ticket price from Equation 14a
- xiv) Calculate $\frac{\partial \sigma_e}{\partial c_p}$ from Equation 12
- xv) Calculate the optimal value of the safety factor from Equation 10 and number of tickets sold
- xvi) If $j = 1$ and $i < 12$ go to viii)
- xvii) If $j = 1$ go to vi)
- xviii) If $i < 24$ go to viii)
- xix) Calculate total airline contribution to profit
- xx) Repeat simulation for different levels of forecasting accuracy, goodwill and elasticity

4. Case Study

We illustrate the method with an example using data drawn from a large European airline which employed an allotments approach with some nesting. We generate a demand function profile as the number of bookings, rather than actual demand, which was not available. We combine the bookings for the two classes of customers into one demand function profile by adding the high class and low class bookings at each time point. There are 23 time points in all, though the periods themselves are not linear, the length of each period decreasing as flight take-off approaches. We initially segment into two stages of the airline booking problem, although this can be extended into as many stages as required by the use of dynamic processing. The first stage, associated with low fare customers, covers the period of time after booking opens while the second stage, associated with high fare customers, covers the later period up to the time of the scheduled flight. We then allocate an initial price for each ticket based on prices charged historically for such seats, so that in our simplified illustration we have one initial price (derived from our profile) for low fare customers and one for high fare customers. Adjustments are made to these initial prices depending on whether there is a surplus or shortfall of remaining capacity as described in section 3.11.

We use airline data on seat bookings to generate a forecast seat demand for each time period for the flight by the use of simulation. Our initial forecast is generated by adjusting the seat demand profile by a random factor so that the forecast is fairly close to the assumed demand with a resultant low

standard deviation of forecasting errors. We then generate another forecast using naïve forecasting (i.e. we simply make our forecast for demand in the next time period equal to the bookings in the current period) with a resultant high standard deviation of forecasting errors. Our final forecast is the weighted average of the two, which provides a range of forecasts from very good to poor depending on the value of the exponential weighting factor, where high values indicate poor forecasting weighted towards the naïve method.

A range of costs describe the aspects of the budget airline yield management problem. The contribution to profit, c_{p_0} is the initial cost of the ticket and is typically determined exogenously through competition in the market place as well as historical flight data. We set the value at 15 units (where each unit is €10) during the first stage (weeks 1-12) and 30 units at the second stage (weeks 13-23). The customer underage (goodwill) cost, c_{u_1} reflects the disappointment cost of not being able to book a seat because there is not one available. In our example we set this at 10% of the ticket cost at stage one ($=0.1 \times 15$) and 20% ($=0.2 \times 30$) at stage two. In the case of the airline underage cost c_{o_2} this reflects the airline's concern at losing a customer booking. This has been set at 20% of the ticket cost at stage one ($=0.2 \times 15$) and 40% at stage 2 ($=0.4 \times 30$). The cost of customer overage, c_{o_1} is the ticket cost reflecting the loss to the customer of booking a seat and not using it. We set this at the cost of the ticket at both stages, but if refunds are possible this would be the difference between the ticket cost and the refund. Unsold tickets reflect the cost of the lost opportunity to the airline, c_{u_2} as loss of contribution to profit, but this may also be the loss incurred when a refund is made. When both of these situations are possible the cost can be calculated as the weighted mean of the ticket cost and the refund cost depending on how often refunds occur compared with non-refunded unsold seats. In our illustration we assume zero refund so that at stage one both the customer overage and airline underage costs are 15 while they are 30 at stage two. Hence $c_{o_1} = 1.5$, $c_{u_1} = 1.5$, $c_{o_2} = 3$, $c_{u_2} = 13.5$. The equilibrium solution with constant variability during stage one is, from Eq. (5)

$$\Phi(k) = \frac{c_{u_1} + c_{o_2} + c_p}{c_{o_1} + c_{u_1} + c_{o_2} + c_{u_2} + c_p} = \frac{1.5 + 3 + 15}{1.5 + 1.5 + 3 + 13.5 + 15} = 0.565 \quad (15)$$

And so $k = 0.164$. For stage two the equilibrium solution gives a value of $k = 0.293$.

The cost settings described in the previous section may not all be available to the airline. We therefore describe another approach based on targets derived from previous experience with allocation obtained from profiling historical demand and pricing information. The first target is that the underage expectation is 9% of mean demand at stage one while it is 6% at stage two. This reflects the ambition to capture more of the customers demanding tickets at stage two than at stage one since the stage two customers pay more for their tickets. This target is derived from the aspiration to capture a certain proportion of demand and depends on the forecasting ability of the airline.

We also derive another measure from the underage target (9% in stage one) using experience and information gathered from earlier flights. This information can be condensed into a measure of

forecasting performance, which is the error standard deviation divided by the mean demand, called the error coefficient of variation (Pearson, 2003). This measure, as our study aims to show, typically differs between the stages, as the airline seats become allocated, enabling the decision maker to set typical underage targets, measured as percentages of the average demand during that time slot. In stage one of the airline data the typical error coefficient of variation (σ_e/μ_D) was 0.279, while at stage two it was 0.223 (Table 2). At stage 1, this reflects high demand and high error variability while there is low demand combined with low error variability at stage 2. This enables the identification of distinct error coefficients of variation at the two stages, though in our example the values turn out to be similar due to the low demand in stage 2. The target % underage at stage one is $u_n = 100\{\phi(k) - k(1 - \Phi(k))\}\sigma_e/\mu_D = 9\%$, which means that $100\{\phi(k) - k(1 - \Phi(k))\} = 9/0.279 = 32.3$, leading to a value of $k=0.164$. The target % overage is $u_n = 100\{\Phi(k) + k\phi(k)\}\sigma_e/\mu_D$. We summarise the above reflections in Table 2, where the numbers highlighted in bold type are derived from market conditions and historical data, while the other values in italics are dependent on these (though the cost coefficients in italics could be evaluated differently rather than as proportions of the ticket cost). These targets for underage and overage will also help later in our study to establish the efficient frontiers in the ‘mix’ phase diagrams illustrated in Figures 4a and 4b.

Table 3 shows the results of applying the solution of Eq. 10 dynamically over the 23 time periods of the booking cycle. The current practice and optimal price appear as c_{p_0} and $c_{p_{opt}}$ respectively. The expected demand and forecast demand at the optimal price appear as μ_D and $\widehat{\mu}_D$ respectively. The expected allocation (supply quantity) at the optimal price appears as μ_Q . μ and η are given by $\mu = \mu_Q - \mu_D$ and $\eta = \mu_Q + \mu_D$, respectively, while σ_e is the standard deviation of the forecasting errors. The forecast and updated remaining seating capacity appear as C_{fst} and C_{new} , respectively. The optimal value of the safety factor, calculated from Eq. 10 is k_{opt} . The value of $c_{p_{opt}}$ is updated in this equation, being increased or decreased depending on whether the value of the updated remaining capacity (C_{new}) is less than or greater than the value of the forecast remaining capacity, C_{fst} . The increase or decrease is calculated using the price elasticity, which for our study is a price elasticity of -1 at stage 1 and -0.5 at stage 2. This reflects the greater sensitivity to price of passengers who book early than those who book late. Eq. (14) is then used to update the forecast demand in line with the price change and elasticity.

	Stage 1	Stage 2	
c_{u_1}	1.5	6	<i>Cost of customer disappointment at unavailable seat or (this cost – Compensation) or weighted average</i>
c_{o_2}	3	12	<i>Cost to airline of concern at loss of customer</i>
c_{p_0}	15	30	<i>Initial ticket cost (often determined by the market)</i>
c_{o_1}	1.5	1.5	<i>Cost to customer of booking and not using seat (such as Ticket Cost – Refund)</i>
c_{u_2}	13.5	28.5	<i>Cost to airline of unsold tickets (such as Ticket Cost or Refund or weighted average)</i>
$\Phi(k)$	0.565	0.615	<i>Steady state newsvendor ratio</i>
k	0.164	0.293	<i>Safety factor</i>
<i>%Underage</i>	9	6	<i>Underage as a % of mean demand</i>
<i>%Overage</i>	13.587	12.537	<i>Overage as a % of mean demand</i>
<i>Error CV</i>	0.279	0.223	<i>Error standard deviation / mean demand</i>
<i>Elasticity</i>	-1	-0.5	<i>Percentage increase in demand for a one percent increase in price</i>

Table 2 Steady State Solution of Eq. (5) at Two Stages

Time	μ_D	μ_Q	μ	η	$\widehat{\mu}_D$	e	σ_e	k_{opt}	c_{p_0}	$c_{p_{opt}}$	C_{fcst}	C_{new}
1	3	3	0	6	3	0		-0.28	15	15	248	247
2	1.3	2	0.7	3.3	2.3	-1		-0.16	15	17.8	246	246
3	5	3	-2	8	3	2	0.71	-0.23	15	15	244	243
4	0.3	2	1.7	2.3	2.3	-2	1.53	-0.06	15	17.8	242	243
5	20.7	18	-2.7	38.7	18.7	2	1.71	-0.39	15	14.42	227	225
6	15.6	20	4.4	35.6	20.6	-5	1.79	-0.32	15	15.95	209	209
7	33	31	-2	64	32	-1	2.95	-0.36	15	15	183	178
8	23.5	26	2.5	49.5	26.5	-3	3.05	-0.28	15	16.75	158	155
9	39.9	36	-3.9	75.9	36.9	3	2.88	-0.33	15	15.81	126	119
10	45.1	42	-3.1	87.1	43.1	2	3.44	-0.30	15	16.53	86	77
11	9.7	12	2.3	21.7	12.7	-3	3.44	-0.13	15	19.97	70	67
12	15.9	15	-0.9	30.9	15.9	0	2.83	-0.26	15	16.75	55	52
13	9.9	11	1.1	20.9	11.9	-2	2.77	-0.33	30	38.4	43	42
14	6.3	8	1.7	14.3	8.3	-2	2.55	-0.16	30	32.8	35	36
15	5.7	7	1.3	12.7	7.7	-2	2.00	-0.16	30	27.2	29	30
16	7.7	8	0.3	15.7	8.7	-1	1.10	-0.25	30	27.2	22	22
17	6	5	-1	11	5	1	0.89	-0.01	30	30	18	17
18	0.3	3	2.7	3.3	3.3	-3	1.30	0.10	30	32.8	15	17
19	6.4	6	-0.4	12.4	6.4	0	1.52	-0.10	30	24.4	11	11
20	4	4	0	8	4	0	1.58	0.06	30	30	8	7
21	3.3	3	-0.3	6.3	2.3	1	1.52	0.17	30	32.8	5	4
22	1.3	2	0.7	3.3	1.3	0	1.64	0.23	30	32.8	3	3
23	2	1	-1	3	1	1	1.52	0.25	30	30	2	2

Table 3. Airline Seat Allocation (Figures shown are for good forecasting)

4.1 Demand Profile

In our method the airline demand data acts as an initial profile over time periods 1-12 (Stage 1) and then over time periods 13-23 (Stage 2). Capacity allocation is then managed by a number of methods. Firstly, the demand profile is required to meet the underage targets (9% and 6% from Table 2) measured against the capacity of the aeroplane. The demand profile is therefore adjusted across all time points to match the underage targets. The demand profile, when adjusted, becomes the capacity plus 1.09 times the initial demand in periods 1 to 12 and the capacity plus 1.06 times the initial demand in periods 13 to 23 as illustrated in Figure 4. The total demand for all 23 weeks is therefore adjusted by scaling to match the capacity of the aeroplane and so is based on a capacitated solution. Secondly, adjustments are made to the ticket price in order to control for the remaining capacity, as described in section 3.11. Thirdly, the remaining capacity is updated at each time period and the optimal seat allocation at the new ticket price is again calculated.

Figure 4 illustrates the demand (Demand) after the above adjustments have been made. It also shows the simulated forecast demand using good quality forecasting, characteristically having a low standard deviation of forecasting errors, of seat demand derived from the adjusted demand data and updated at each time point during the two stages. Finally, it shows the profile of poor quality (naïve) forecasting, having a high standard deviation of forecasting errors, of seat demand derived from the naïve forecasting method which uses the latest demand value as the forecast for the demand in the next time period.

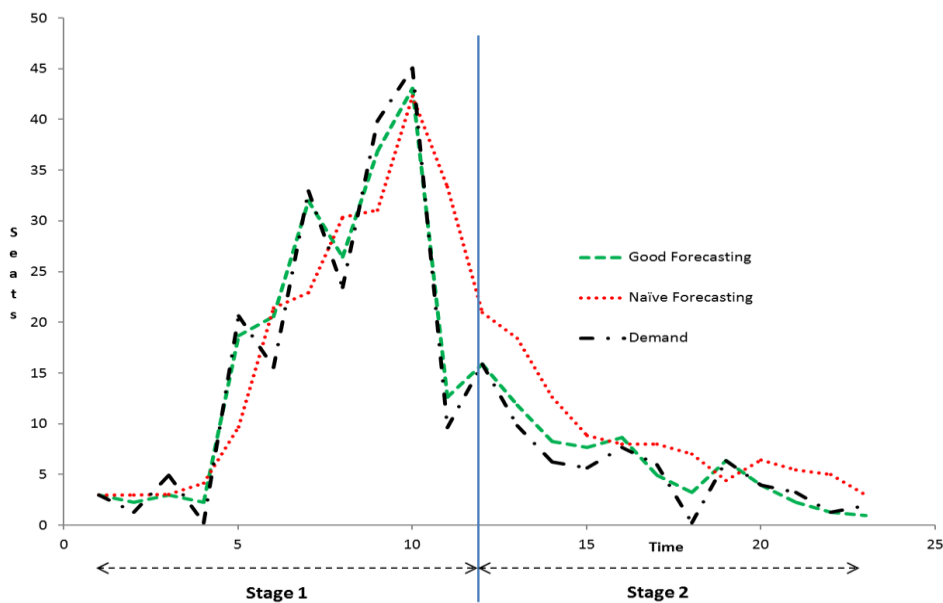


Figure 4. Demand Profile with Good and Naïve Forecasting

In subsequent analysis a weighted average of forecast demand (where α is the weighting parameter with value zero for good quality forecasting and value 1 for naïve forecasting) from the two methods is used to generate a new forecast. The factor levels for Expsmooth, Goodwill and Elasticity at the

two stages are shown in Table 4. For the simulation, Expsmooth ranges from 0 ($\alpha=0$) to 2($\alpha=0.6$) so that the value of Expsmooth corresponding to $\alpha=1$ in Figure 4 would be 3.33.

Stage 1	Expsmooth	Goodwill	Elasticity
First level	0 when $\alpha = 0.0$	0 when $c_{u_1} = 0, c_{o_2} = 0$	1 when <i>Elasticity</i> = -1
Second level	1 when $\alpha = 0.3$	1 when $c_{u_1} = 1.5, c_{o_2} = 3$	2 when <i>Elasticity</i> = -2
Third level	2 when $\alpha = 0.6$	2 when $c_{u_1} = 3, c_{o_2} = 6$	3 when <i>Elasticity</i> = -3

Stage 2	Expsmooth	Goodwill	Elasticity
First level	0 when $\alpha = 0.0$	0 when $c_{u_1} = 0, c_{o_2} = 0$	1 when <i>Elasticity</i> = -0.5
Second level	1 when $\alpha = 0.3$	1 when $c_{u_1} = 6, c_{o_2} = 12$	2 when <i>Elasticity</i> = -1.0
Third level	2 when $\alpha = 0.6$	2 when $c_{u_1} = 12, c_{o_2} = 24$	3 when <i>Elasticity</i> = -1.5

Table 4. Factor Levels in the Experimental Design

Figure 5a shows the mix and global phase diagrams for the two stages of the pricing procedure when the forecasting method is good (expsmooth = 0), goodwill is average (goodwill = 1) and elasticity is low (elasticity level = 1) representing typically stable conditions commonly associating with ‘lean’ systems. In the mix diagram for stage 1, variability and expected volume (demand + supply) are higher than at stage 2. Progress during this stage is mapped by the mix plot and shows good control, contained as it is within the efficient frontiers. The target solution (1.21, 6.07) shown in the figure is found by substituting the values for k ($= 0.164$ derived from the underage target of 9%), c_{p_0} ($=15$) and the error CV ($=0.279$ in Table 2) into Eq. (1). This is explained in Appendix B which describes the way in which the efficient frontiers are mapped. The plot oscillates between the overage (surplus capacity) and underage (demand exceeds capacity) efficient frontier boundaries reflecting the pricing adjustments occurring in Table 3. So, for instance, actual remaining capacity at the end of time point 1 is 247 ($C_{new}=247$) since three tickets were sold ($\mu_Q =3$). Since the remaining capacity is less than the forecast capacity ($247=C_{new}<C_{fcst}=248$) the price at time point 2 increases from the base price ($17.8 =c_{p_{opt}}> c_{p_0}=15$). The amount of the increase is calculated using the price elasticity and compensates for the shortage in capacity by reducing demand because of the higher price. This has the effect of reducing the expected demand ($\mu_D =1.3$) in time point 2 which results in a levelling of capacity at the end of this time point ($246=C_{new}=C_{fcst}=246$).

Phase Plane Plots for Good Quality Forecasting - Lean (0,1,1) Profit = 4726.40

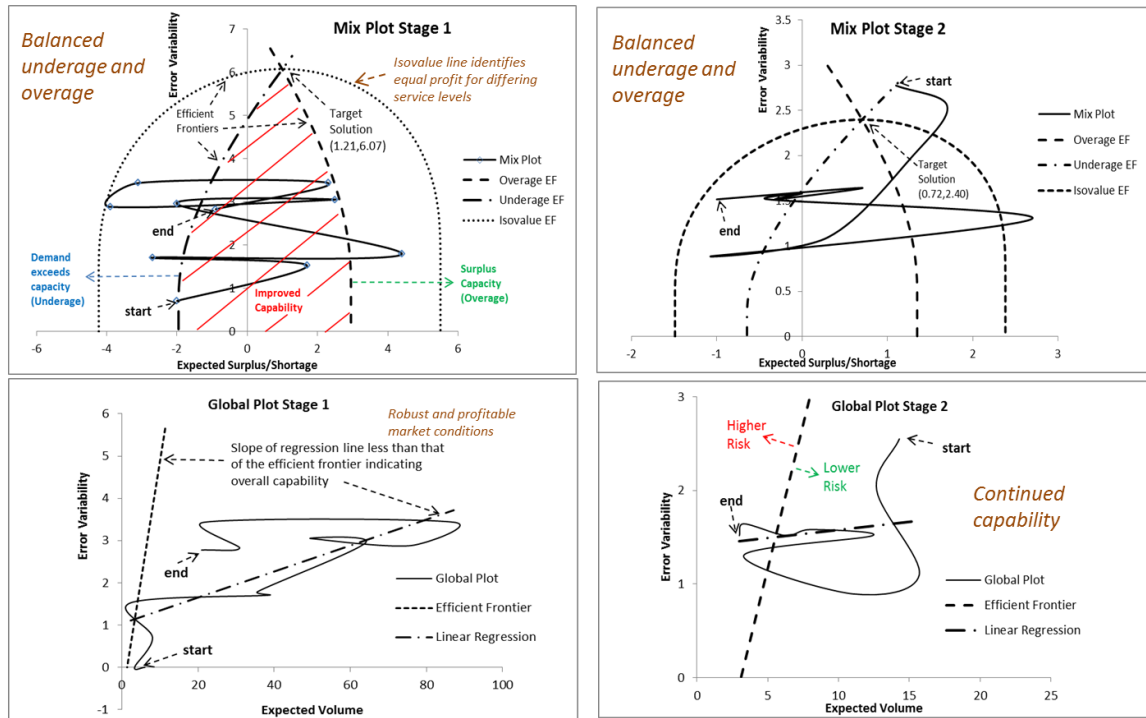


Figure 5a. Phase Diagrams for Good Quality Forecasting, Average Goodwill and Low Elasticity

Figure 5a also illustrates the mix diagram for stage 2 where variability and expected volume (demand + supply) are lower than at stage 1. The mix plot again shows good control, contained as it is within the efficient frontiers. The target solution (0.72, 2.40) shown in the figure is found (again, similar to Appendix B) by substituting the values for k ($= 0.293$ derived from the underage target of 6%), c_{p_0} ($=30$) and the error CV ($=0.223$ in Table 2) into Eq. (1). The plot again oscillates between the overage (surplus capacity) and underage (demand exceeds capacity) efficient frontier boundaries reflecting the pricing adjustments occurring in Table 3. So, for instance, actual remaining capacity at the end of time point 14 is 36 ($C_{new}=36$) since seven tickets were sold ($\mu_D =7$). Since the remaining capacity is greater than the forecast capacity ($36=C_{new}>C_{fcst}=35$) the price at time point 15 decreases from the base price ($27.2=c_{p_{opt}}>c_{p_0}=30$). The amount of the decrease is again calculated using the price elasticity and compensates for the shortage in capacity by reducing demand because of the higher price. This has the effect of increasing the expected demand ($\mu_D =5.7$) in time point 15 which results in a levelling of capacity but not until the end of time point 16 following a repeated price reduction ($22=C_{new}=C_{fcst}=22$).

Finally Figure 5a illustrates the global phase diagrams for stages 1 and 2. In both of these phase diagrams the slope of the regression line (plotted using the points on the path) is less than the slope of the efficient frontier derived from Eq. 9. This shows that the allocation method is coping robustly with the simulated market conditions.

Phase Plane Plots for Poor Forecasting and High Elasticity (2,1,3) Profit = 4405.98

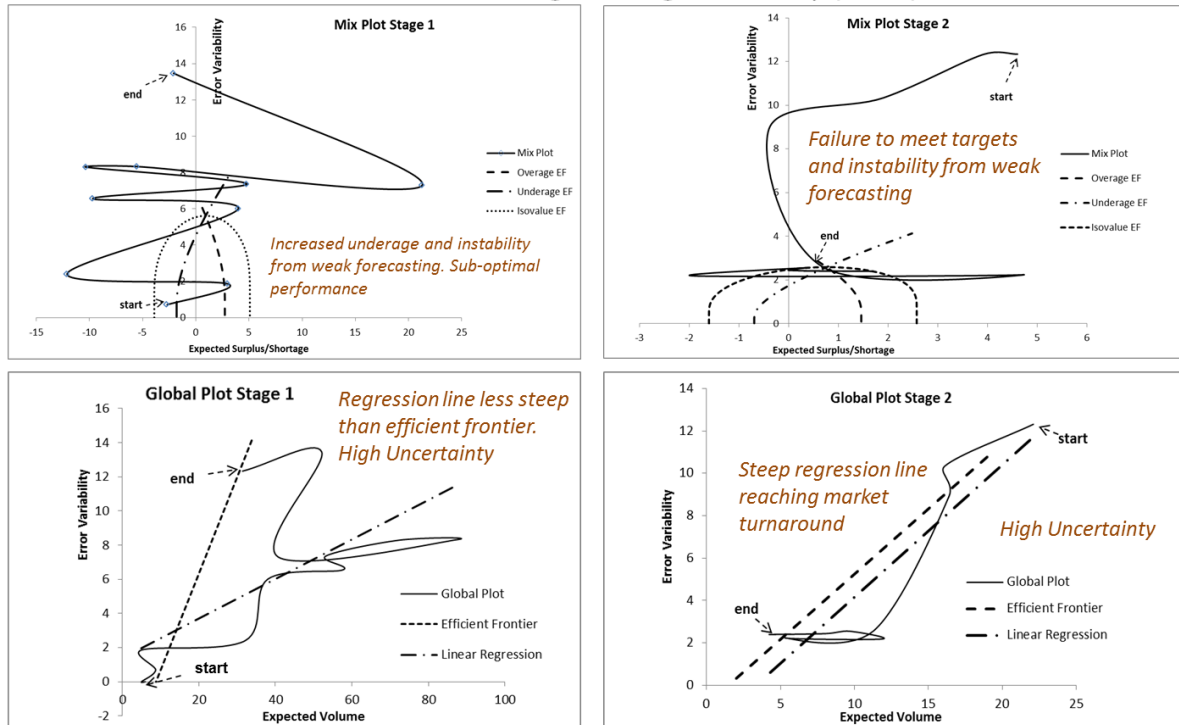


Figure 5b. Phase Diagrams for Poor Quality Forecasting, Average Goodwill and High Elasticity

In contrast to the stable conditions illustrated in Figure 5a, Figure 5b shows the mix and global phase diagrams for the two stages of the pricing procedure when the forecasting method is poor ($\text{expsmooth} = 2$), goodwill is average ($\text{goodwill} = 1$) and elasticity is high ($\text{elasticity} = 3$) representing typically unstable conditions commonly associating with ‘agile’ systems. The mix plot for stage 1 illustrates increased underage due to the way in which naïve forecasting lags behind increasing demand. The mix plot for stage 2 also illustrates the suboptimal performance and the failure to meet targets. The stage 2 global plot has a regression line which reaches the slope of the efficient frontier indicating a turning point in allocation capability in the presence of high uncertainty. Both types of plot indicate the need for a reassessment of the process and a setting of more appropriate targets. They might also assist in the decision to cancel future flights with recurring profile patterns.

4.2. Sensitivity Analysis and Diagnostics

We carried out a series of tests based on a full factorial design (Table 4) with response variable ‘Contribution to Profit’ and factors Forecast Quality (expsmooth with levels 0,1 and 2), Goodwill (with levels 0,1,2) and Elasticity (levels 1,2,3).

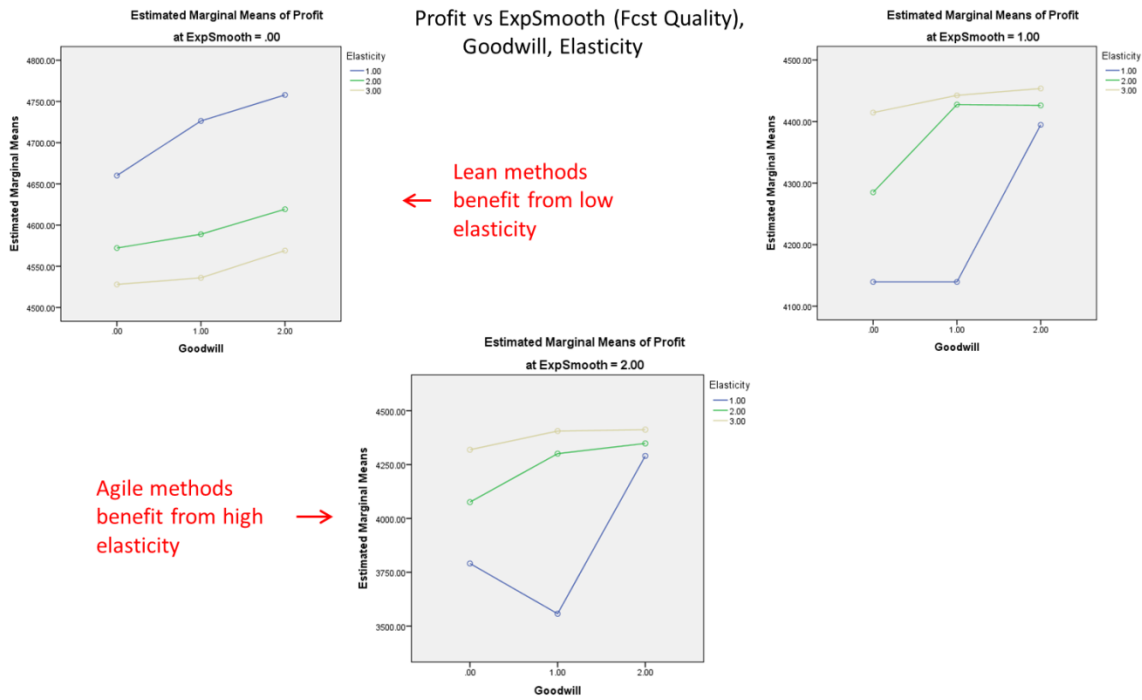
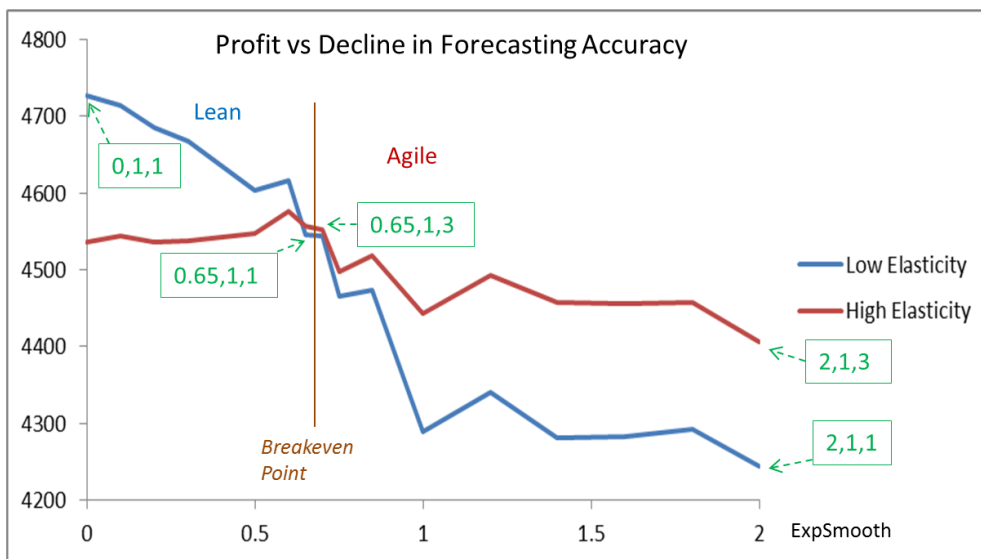


Figure 6. Analysis of Means Test for Full Factorial Design (Airline Profit)

An analysis of means test produced the results shown in Figure 6. The profit used is the airline profit calculated as the revenue from ticket sales. This illustrates the way in which lean methods, typically associated with good forecasting and stable conditions, benefit from low elasticity and where goodwill is a key factor in profitability. On the other hand it shows that agile methods, typically associated with high variability and unstable conditions, benefit from high elasticity. It also demonstrates that incorporating goodwill is generally profitable but moderated by elasticity.

The transition from LEAN to AGILE

Goodwill = 1 (Fixed at centre) ; Elasticity = 1 (Low) ; 3 (High)



The transition occurs rapidly between the values of 0.5 and 1

Figure 7. Interaction Plot illustrating Transition from Lean to Agile (Airline Profit)

Figure 7 is an interaction plot illustrating the transition from lean to agile methods, where we again use the airline profit as dependent variable. In the case of our airline example we use the term ‘lean’ to represent methods associated with stable demand while the term ‘agile’ denotes unstable demand sometimes associated with budget airlines. At the breakeven point we see that elastic and inelastic demand have equal profitability, with the transition occurring rapidly between the values of 0.5 and 1 for the exponential smoothing parameter. The plot illustrates the way in which increased profitability is associated with low elasticity in lean methods while being associated with high elasticity for agile methods.

Phase plane Plots at Transition between Lean and Agile (0.65,1,1) Profit = 4545.33

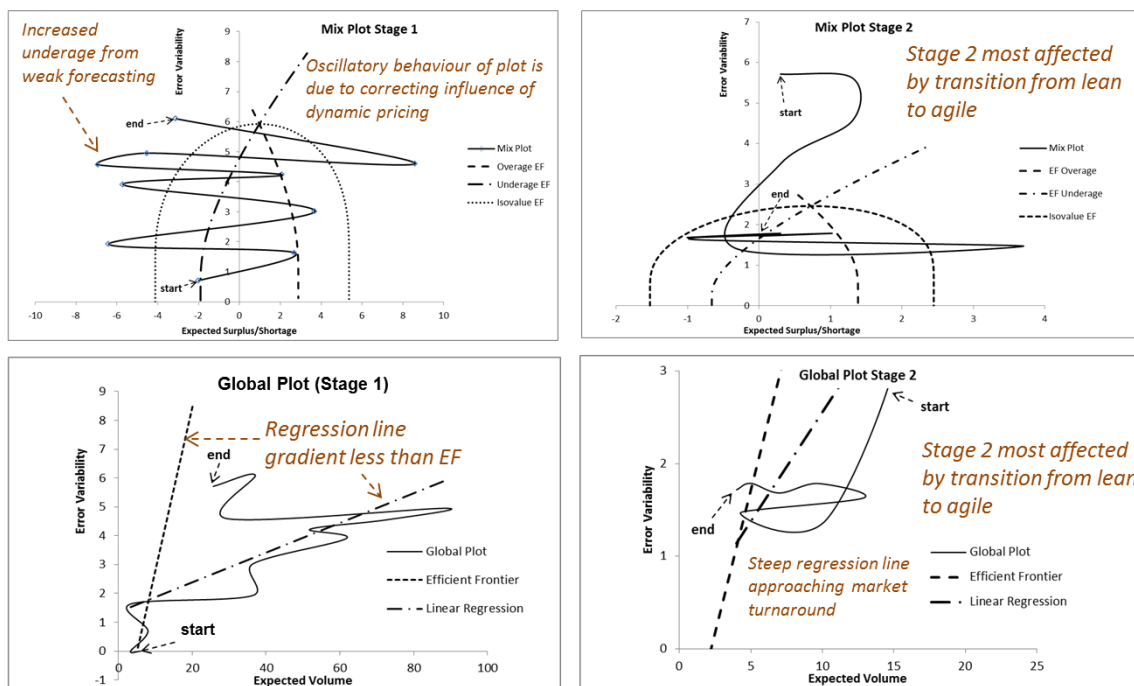


Figure 8. Phase Diagrams at the Transition between Lean and Agile

Figure 8 illustrates the phase diagrams at the breakeven transition between lean and agile methods where the profit levels are similar. We illustrate the phase diagrams for the settings (0.65,1,1) for low elasticity but these are very similar to the phase diagrams for the settings (0.65,1,3) where the elasticity is high. The mix plot for phase 1 illustrates the oscillatory behaviour and more extreme values as the forecasting quality deteriorates while at stage 2 the plot starts in a region well outside of the target area. The global plots show a movement from fairly stable operation in stage one to a less stable position in stage two where the regression plot approaches the maximum slope for optimal operation. These plots are of key interest because, in spite of a reduction in forecast quality, the system is comparatively insensitive to the price elasticity of demand. Therefore, if forecasting quality can be maintained at around this level, there will be greater robustness in the system.

In the final stage of our analysis we repeated the full factorial experiment making use of the details from 18 flights drawn from the airline data set. The results were very similar to those described above and in Figures 6, 7 and 8, though the breakeven point described in Figure 7 varied between the flights. In fourteen of the cases the value of the exponential smoothing parameter, α , at the breakeven point lay between 0.3 and 0.8 (the factorial design was later repeated for values of α ranging between 0 and 1). There were four outliers, three with high values close to 2 while one value was low at 0.1. The high values were characterised by smooth transitions in the pattern of the demand so that naïve forecasting still gave good results since the demand from one time point to the next was very smooth. The low value was characterised by the reverse feature so that the transitions were not at all smooth. As a final test we varied the value of r_σ (which controls the rate at which the quadratic expression for variability changes with respect to price) around its initial value of 0.1 and found little sensitivity of the model in this area, though for values of r_σ greater than 0.5 the model restricted the optimal solution much closer to current practice policy, as would be expected.

5. Conclusion

We develop an effective tool to monitor and map the shape of uncertainty in order to better address a number of issues raised in revenue management. Among these issues is the distinction between exogenous influences from global factors such as competition in the market place and endogenous influences from local factors such as a healthy relationship between the customer base and the airline operator. For this purpose we develop ‘mix’ and ‘global’ phase diagrams which map and monitor variability against the endogenous and exogenous variables. Another issue is the distribution-free methodology which does not make assumptions about the demand distribution but only about the distribution of the forecasting errors which makes it adaptable to most demand management systems. A further issue is that the methodology can be effective in identifying when the process is in control since it maps the progress of the allocation process from a local and a global perspective incorporating efficient frontiers to monitor when it is sub-optimal or out of control. The model takes into account limitations in the performance of the forecaster preventing significant drifting from the area where the forecaster is most reliable. We address the issue of goodwill costs which are included in the model and form a key role in uncovering optimal profit levels in a market influenced by demand elasticity. Finally the staged approach allows for discerning division of the demand profile so that fare classes can be allocated through flexible adjustments made at the point of sale. Furthermore, each stage is chosen with the help of two significantly uncorrelated phase diagrams and the results used to make improved choice of stages in subsequent allocations.

Among our findings is the demonstration that lean methods, defined in the context of airline revenue management (but with broader applications in inventory and other systems), benefit from low demand elasticity through increased profit levels. Furthermore that the incorporation of customer

goodwill costs into the model is generally profitable but less so when there is high demand elasticity. We also found that when the demand profile is staged it is the later stage customers that are most affected by the transition from stable to less stable conditions. Another finding is that the dynamic pricing model produced oscillations in the ‘mix’ phase diagrams due to the correcting effect of the price changes which bring the system back to its target performance measures. Finally we find that the breakeven point (measured by forecasting quality) defines the equilibrium between lean and agile methods with regard to elasticity when profit is the response variable.

Although the methods described here apply to the single leg revenue management problem they can be extended for use generally across revenue management as well as in other contexts. The allocation of rooms in a large hotel would benefit from these methods since the capacity of the hotel, like that of an aeroplane, is fixed. Another area for extension is the network model where a common approach has been the use of ‘bid price’ methods as well as improvements to such methods (Meissner et al., 2012). The methods can also be extended to investigate the dynamics of cancellations and overbooking (Gosavi et al., 2007) by a relaxation of the steady state capacity requirement that $C = \mu_Q$. Inspection of the phase diagrams can be used to assess forecasting methods and customer satisfaction, as well as monitor progress and decide on whether or not to accept a profiled base price and allocation strategy.

There are some limitations associated with some of the outlined methods. The first is that bias can enter into forecasting methods which results in a violation of the normality of the forecasting errors assumption. Care needs to be taken when interpreting the phase diagrams in this event using them mainly as indicators of when the process becomes out of control. Another limitation is when the forecasting errors are not normally distributed which often occurs for small sample sizes. If, for instance, only a few seats are booked in a period then the distribution is more likely to be Poisson. This, however, can be dealt with by grouping together two or more periods of data with some resultant loss in allocation precision.

References

- Alfares, H. K., Elmorra, H. H. (2005) The distribution-free newsboy problem: Extensions to the shortage penalty case. *Int. J. Prod. Econ.* 93-94, 465-477
- Belobaba, P. P., (1989) Application of a Probabilistic Decision Model to Airline Seat Inventory Control, *Operations Research*, 37, (2), 183-197
- Belobaba, P.P, Odoni, A., Barnhart, C. (Eds.) (2009) The Global Airline Industry. J.Wiley & Sons.
- Birbil, S.I., Frenk, J.B, Gromicho, J.A., Zhang, S. (2009) The role of robust optimization in single-leg airline revenue management. *Management Science*, 55, (1), 148-163
- Dana, J. D. (2008) “New Directions in Revenue Management Research,” *Production and Operations Management*, Vol. 17, (4), 399-401.

- Elmaghraby, W. and Keskinocak, P. (2003). Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science*, 49, (10), 1287–1309.
- Gallego, G. and Sahin, O. (2010). Revenue Management with Partially Refundable Fares. *Operations Research*, 58, (4), 817-833
- Gosavi, A., Ozkaya, E. and Kahraman, F. (2007) Simulation optimization for revenue management of airlines with cancellations and overbookings. *OR Spectrum*, 29 (1) 21-38
- Kocabıyıkoglu, A., Popescu, I. and Stefanescu, C. (2014) Pricing and Revenue Management: The Value of Coordination. *Management Science*, 60, (3), 730-752
- Lee, A.O., (1990). Airline reservations forecasting: probabilistic and statistical models of the booking process. PhD dissertation, MIT Flight Transportation Laboratory, Cambridge, MA.
- Levin, Y, McGill, J and Nediak, M. 2008. Risk in Revenue Management and Dynamic Pricing. *Operations Research*, 56, (2), 326-343
- Levin, Y, McGill, J (2009) Introduction to the special issue on revenue management and dynamic pricing *European Journal of Operational Research* 197 (2009) 845–847
- Littlewood, K. (1972) “Forecasting and Control of Passenger Bookings,” in *AGIFORS Symposium Proc.* **12**, Nathanya, Israel.
- Marcus, B. and Anderson, C. (2008) “Revenue Management for Low Cost Providers”, *European Journal of Operational Research* 188, 258-272
- McGill, I.M., Van Ryzin, G.J. (1999) Revenue Management: Research Overview and Prospects, *Transportation Science* 33, (2), 233-256
- Meissner, J., Strauss, A., (2012) Improved bid prices for choice-based network revenue management. *European Journal of Operational Research* 217, (2), 417–427
- Ming, S.X. and Zhang, F. (2009) “On the Value of Commitment and Availability Guarantees when Selling to Strategic Consumers,” *Management Science* 55, No. 5, 713-726.
- Modarres, M., Sharifyazdi, M. (2009) Revenue management approach to stochastic capacity allocation problem. *European Journal of Operational Research* 192, (2), 442-459
- Ozkan, C., Karaesmen, F., Ozekici, S. (2014) A revenue management problem with a choice model of consumer behaviour in a random environment. *OR Spectrum*. DOI 10.1007/s00291-014-0380-9
- Pasternack, B. (1985) Optimal pricing and return policies for perishable commodities. *Marketing Science*, 4, (2), 166–176.
- Pearson M. (2003). “An Equilibrium Solution to Supply Chain Synchronization” *IMA Journal of Management Mathematics*, 14, (3), 165-185.
- Pearson M., Masson R. and Swain A. (2010) ‘Process Control in an Agile Supply Chain Network’ *International Journal of Production Economics*, 128, 22-30.
- Phillips, R. L. (2005) Pricing and Revenue Optimization. Stanford University Press.

- Poelt, S. (1998) “Forecasting is difficult – especially if it refers to the future” *Reservations and Yield Management Study Group Annual Meeting Proceedings, Melbourne, Australia, AGIFORS*
- Raza, S. A. (2014) A distribution free approach to newsvendor problem with pricing *4OR-Q J Oper Res* 12:335-358
- Sahin, E., Dallery, Y., (2009) Assessing the impact of inventory inaccuracies within a Newsvendor framework. *European Journal of Operational Research* 197, (3), 1108–1118.
- Scarf H. (1958) A min-max solution of an inventory problem. In: Arrow K, Karlin S and Scarf H (eds). *Studies in the Mathematical theory of Inventory and production*. Stanford University Press: California, pp 201–209.
- Talluri, K. T. and Van Ryzin, G.J. (2004) *The Theory and Practice of Revenue Management*, volume 68 of International Series in Operations Research and Management Science. Springer.
- Weatherford, L. R., Belobaba, P. P., (2002) Revenue impacts of fare input and demand forecast accuracy in airline yield management. *Journal of the Operational Research Society* 53, 811-821
- Weatherford, L. R., Bodily, S. E., (1992) A Taxonomy and Research Overview of Perishable-Asset Revenue Management, Overbooking, and Pricing. *Operations Research* 40, (5), 831-844
- Vulcano, G., Van Ryzin, G. and Maglaras, C. (2002) Optimal Dynamic Auctions for Revenue Management. *Management Science* 48, (11), 1388-1407

Appendices

Appendix A

Changing Variability

The primal-dual objective function makes use of the primal-dual transformation. We consider the mix variable, which is the difference between the dual (Q) and primal (D) variables. We define $X=Q-D$ and forecast values $\hat{X} = \hat{Q} - \hat{D}$. We let σ_e be the standard deviation of $e = e_Q - e_D$, where e_D and e_Q are the demand and supply forecasting error terms, respectively. Then $\hat{X} - X + k\sigma_e \sim N(k\sigma_e, \sigma_e^2)$, and $E(X) = \mu = \mu_Q - \mu_D$. e will have a univariate normal distribution with mean zero. A similar result follows for the global variable, η . The transformation applied is: $\mu = \mu_Q - \mu_D$ and $\eta = \mu_Q + \mu_D$, so that $\mu_D = (\eta - \mu)/2$, $\mu_Q = (\eta + \mu)/2$, and $\partial\mu_D/\partial\mu = -1/2$, $\partial\mu_Q/\partial\mu = +1/2$, $\partial\mu_D/\partial\eta = +1/2$ and $\partial\mu_Q/\partial\eta = +1/2$. We are then able to find the maximum value of the objective function (1) subject to the constraint (2) by applying a Lagrange multiplier with decision variables μ , η and k . We have

$$L = \mu_D c_p - \{(\phi(k) - k(1 - \Phi(k)))c_p + (k\Phi(k) + \phi(k))(c_{o_1} + c_{u_2}) + (\phi(k) - k(1 - \Phi(k)))(c_{u_1} + c_{o_2})\}\sigma_e - \lambda(\mu - k\sigma_e).$$

So that $\frac{\partial L}{\partial k} = -\sigma_e \{\Phi(k)(c_{o_1} + c_{u_2} + c_{u_1} + c_{o_2} + c_p) - (c_{u_1} + c_{o_2} + c_p)\} + \lambda\sigma_e = 0$,

and $\lambda = \Phi(k)(c_{o_1} + c_{u_2} + c_{u_1} + c_{o_2} + c_p) - (c_{u_1} + c_{o_2} + c_p)$. Also

$$\frac{\partial L}{\partial \mu} = -\frac{1}{2}c_p - \sigma_e^\mu \left\{ (\phi(k) + k\Phi(k))(c_{o_1} + c_{u_2} + c_{u_1} + c_{o_2} + c_p) - k(c_{u_1} + c_{o_2} + c_p) \right\} - \lambda(1 - k\sigma_e^\mu) = 0,$$

where $\sigma_e^\mu = \frac{\partial \sigma_e}{\partial \mu}$ so that $\lambda = -\frac{1}{2}c_p - \sigma_e^\mu [\phi(k)(c_{o_1} + c_{u_2} + c_{u_1} + c_{o_2} + c_p)]$
 $= \Phi(k)(c_{o_1} + c_{u_2} + c_{u_1} + c_{o_2} + c_p) - (c_{u_1} + c_{o_2} + c_p).$

Eq. (3) follows. Furthermore

$$\frac{\partial L}{\partial \eta} = \frac{1}{2}c_p - \sigma_e^\eta \left\{ (\phi(k) + k\Phi(k))(c_{o_1} + c_{u_2} + c_{u_1} + c_{o_2} + c_p) - k(c_{u_1} + c_{o_2} + c_p) \right\} + \lambda k \sigma_e^\eta = 0,$$

where $\sigma_e^\eta = \frac{\partial \sigma_e}{\partial \eta}$, which leads to Eq. (4).

The Hessian determinant is negative definite for $\sigma_e^{\mu\mu}\sigma_e^{\eta\eta} - (\sigma_e^{\mu\eta})^2 > 0$. If there were no clear maximum, then the structure of the Hessian would need to be investigated. It could be that the profit increases without limit as the global output increases. This would happen if $\sigma_e^\eta < 0$, for instance.

Incorporating Price Elasticity

We again use a Lagrange multiplier, but this time we take into account the dependence of the ticket price, c_p on the global variable, η . Eqs. (8) and (9) follow. We see that Eq. (3) is the same as Eq. (8) which reflects the assumption that c_p does not depend on μ . The difference between Eqs. (4) and (9) reflects the dependence of c_p on η .

Appendix B (Equation of Isovalue Line)

This is derived by substituting the optimal values for the overage (o_v) and underage (u_n) into Eq. 1:

E(Profit) =

$$\begin{aligned} \mu_D c_p - \{ (\phi(k) - k(1 - \Phi(k)))c_p + (k\Phi(k) + \phi(k))(c_{o_1} + c_{u_2}) + (\phi(k) - k(1 - \Phi(k)))(c_{u_1} + c_{o_2}) \} \sigma_e \\ = \mu_D c_p - \{ (\phi(k) - k(1 - \Phi(k)))(c_{u_1} + c_{o_2} + c_p) + (k\Phi(k) + \phi(k))(c_{o_1} + c_{u_2}) \} \sigma_e \\ = \mu_D c_p - \{ (c_{u_1} + c_{o_2} + c_p)u_n + (c_{o_1} + c_{u_2})o_v \} \end{aligned}$$

And so $\mu_D c_p - E(\text{profit}) = \{ (c_{u_1} + c_{o_2} + c_p)u_n + (c_{o_1} + c_{u_2})o_v \}$ (B.1)

On the isovalue line the LHS of Eqs. (1) and (B.1) are constant. Setting $k = +\infty$ in Eq. (1),

$$\begin{aligned} \mu_D c_p - E(\text{profit}) = \\ \{ (\phi(k) - k(1 - \Phi(k)))c_p + (k\Phi(k) + \phi(k))(c_{o_1} + c_{u_2}) + (\phi(k) - k(1 - \Phi(k)))(c_{u_1} + c_{o_2}) \} \sigma_e \\ = k\sigma_e \{ (\phi(k)/k - (1 - \Phi(k)))c_p + (\Phi(k) + \phi(k)/k)(c_{o_1} + c_{u_2}) + (\phi(k)/k - (1 - \Phi(k)))(c_{u_1} + c_{o_2}) \} \\ = \mu \{ (\phi(k)/k - (1 - \Phi(k)))c_p + (\Phi(k) + \phi(k)/k)(c_{o_1} + c_{u_2}) + (\phi(k)/k - (1 - \Phi(k)))(c_{u_1} + c_{o_2}) \} \\ = \mu(c_{u_1} + c_{o_2}) \end{aligned}$$
 (B.2)

Equating the RHS of Eq. (B.1) and Eq. (B.2) gives $\mu = [o_v + ((c_{u_1} + c_{o_2} + c_p)/(c_{o_1} + c_{u_2}))u_n]$. The

second result follows by setting $k = -\infty$ in Eq. (1).